

The evolution of build-up package technology and its design challenges

This paper reviews sequential build-up (SBU) laminate substrate development from its beginning in 1988. It reports on developments in this technology for IBM applications since its adoption in 2000. These laminated substrates are nonuniform structures composed of three elements: a core, build-up layers, and finishing layers. Each element has evolved to meet the demands of packaging applications. Thin-film processing has greatly enhanced the wiring capability of SBU laminate substrates and has made this technology very suitable for high-performance designs. This paper focuses on the challenges encountered by IBM during the design, manufacture, and reliability testing phases of development of SBU substrates as solutions for application-specific integrated circuit (ASIC) and microprocessor packaging applications.

E. D. Blackshear
M. Cases
E. Klink
S. R. Engle
R. S. Malfatt
D. N. de Araujo
S. Oggioni
L. D. LaCroix
J. A. Wakil
G. G. Hougham
N. H. Pham
D. J. Russell

Introduction

The increasing demand for computer performance has led to higher chip internal clock frequencies and parallelism, and has increased the need for higher bandwidth and lower latencies. Processor frequencies are predicted to reach 29 GHz by 2018, and off-chip signaling interface speeds are expected to exceed 56 Gb/s [1, 2]. Optimization of bandwidth, power, pin count, or number of wires and cost are the goals for high-speed interconnect design. The electrical performance of interconnects is restricted by noise and timing limitations of the silicon, package, board and cable.

As a result of rapidly emerging technologies and applications, the boundaries between semiconductor, packaging, and system technologies are no longer clear; they must all be considered concurrently in a system-level approach in order to optimize the substrate design. There is an increased awareness in the semiconductor industry that assembly and packaging is an essential and integral part of the semiconductor product. Packaging technology has become a critical competitive factor in many market segments, since it affects operating frequency, power, reliability, and cost.

Sequential build-up (SBU) laminate substrate technology is now the technology of choice for high-density, high-performance silicon packaging. In 1997, SBU technology was selected by Intel [3] for flip-

chip packaging and has been widely adapted for this application. This paper reviews the invention of laminate substrate packaging and discusses rapidly evolving trends in the evolution of SBU, including its broadening application in IBM servers for high-speed system-level interconnects. The importance of properly designing the substrate for high-speed signaling is discussed, including identification of key parameters and design tradeoffs for both application-specific integrated circuit (ASIC) and microprocessor chip designs.

Applications place performance demands on packaging which can best be met by organic materials. The key attributes of organic laminate package technologies as they pertain to the electrical performance of the subassembly are highly electrically conductive metallurgy to minimize resistive voltage drops and to effectively deliver power to the chip; low-inductance connections to reduce simultaneous switching noise; low-dielectric-constant insulator materials to better match board impedances and to reduce undesirable parasitic capacitances; and advanced thermal interface materials to manage high power densities on the chip and to improve performance. The importance of properly designing the substrate for these applications is emphasized, including a discussion of the identification and control of key physical design parameters and a description of the design optimization technique.

©Copyright 2005 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/05/\$5.00 © 2005 IBM

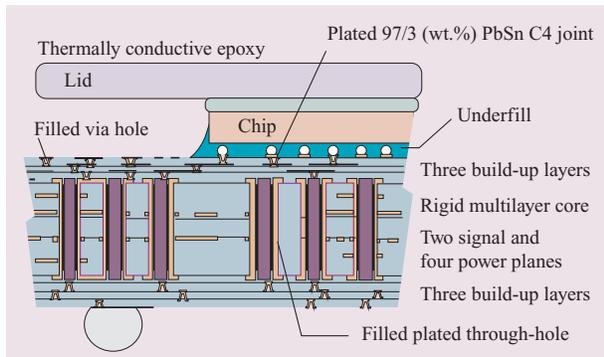


Figure 1

Typical cross section of a flip-chip ball grid array.

History of technology development

Technology origins

Wire-bond interconnects have been the workhorse technology [4] for industry microprocessors and their associated supporting chips since the inception of the personal computer industry. Its primary advantages have been low cost, design flexibility, and thoroughly demonstrated reliability. Its major limitation is wiring capability in terms of both total numbers of signals and electrical performance.

Flip-chip interconnection has been a core IBM approach to silicon packaging [5] for much of the history of IBM servers. It provides the highest interconnect density from a chip to redistribution circuitry, or substrate, that is currently possible. Until recently, the interconnect density enabled by flip-chip technology could be provided only by multilayer ceramic substrates, which have a manufacturing cost significantly higher than that of other packaging circuit elements of a system, such as printed circuit boards (PCBs). Analysis of published design ground rules shows that build-up laminate substrate technology is a breakthrough approach to flip-chip interconnection, achieving significantly higher per-layer wiring densities than printed circuit boards that use many of the same basic materials and processes. Build-up laminate cost is significantly less than that of ceramic dielectric-based competitive technologies. It offers electrical performance enhancements as well through the use of copper conductors and lower-dielectric-constant insulator materials.

The concepts of SBU lamination and its enabling technologies, such as semi-additive plating and high-volume, low-cost micro-via formation, have been under discussion at the leading edge of the PCB industry for some time [6]. The technique which eventually developed into today's substrate industry was pioneered at the IBM

Yasu facility in Japan. In this approach, called surface laminar circuitry (SLC) [7], Probimer**, a material originally developed as a solder resist, was used as an interlayer dielectric. Photoimaging processes were used for via formation. Competitive but significantly more expensive via formation processes at that time were mechanical drilling, plasma etching, and laser drilling. SLC was selected by IBM in 1988 as the sole technology for direct attachment of silicon chips to circuit boards. A token ring adapter card based on SLC began production in 1992, and SLC was subsequently adapted for Personal Computer Memory Card International Association (PCMCIA) format and video adapter cards.

IBM explored the application of SBU to single flip-chip modules but found no high-value applications. When flip-chip technology was originally developed by IBM for ceramic applications, the engineers selected 97/3% lead/tin (wt.%) as its solder bump material. To reflow for bonding, this material requires a temperature of more than 310°C, which would irreversibly degrade most epoxy-based organic substrates. The use of lower-melting-point solders with higher tin content for flip-chip modules on SBU was found to put increased demands on underfill (**Figure 1**), an encapsulating resin used to reinforce the post assembly of C4 (controlled collapse chip connection) joints. It also affected solder fatigue performance. As a solution, eutectic solder paste was adapted [7] and applied to the substrate for joining high-lead-content chip bumps to organic substrates. This provided the fatigue properties of high-lead-content balls with the low melting point of high-tin solder. Owing to the reduced volume of high-tin solder, the probability of underfill shorts was reduced.

As the level of silicon integration in PC processors increased, the I/O density of flip-chip packaging was needed to accommodate higher pin counts and operating frequencies. Intel [3] developed the first high-volume applications of flip-chip technology using SBU packaging. After initial discussions with IBM Yasu, they partnered with Ibiden Corporation to develop a next-generation clone of the SLC process. The primary advance was the development of a new dielectric material, a hybrid copolymer of epoxy thermoset and linear thermoplastic material, which provided an interpenetrating polymer network (IPN) [8] that offered enhanced mechanical properties. The composite substrate resulting from this collaboration was called the IPN build-up substrate system (IBSS) [9]. It was introduced by Intel in 1997 as a package for the Pentium 2** processor. This offering included high-density flip-chip ball-grid array (FC-BGA) packages for mobile applications (**Figure 1**) and flip-chip ball grid array on a pinned interposer for socketed desktop and server applications. Marketed by Intel as an organic land grid array (OLGA),

the new package offered reduced electrical noise and accommodated high pin counts at a considerably lower cost [3, 10], with substantial improvements in power distribution and signal transmission characteristics due in part to the low dielectric constant of the build-up material.

Manufacturing process

In conventional printed circuit board (PCB) processing, circuitry is formed in a largely parallel process on individual cores, with composite layers consisting of a fully cured fiberglass-reinforced epoxy dielectric clad with copper foil. Cores are formed into a multilayer circuit board by interleaving with layers of partially cured, unclad dielectric material and joining under heat and pressure into a composite in a single lamination step. Layer-to-layer interconnection is achieved by mechanically drilling holes through the full thickness of the composite and depositing copper plating on the hole walls, joining inner-layer circuitry where it intersects the plated barrel. The total wiring density achievable through this conventional process is limited by etching tolerances and dielectric thicknesses required to achieve a target impedance, e.g., 50 Ω.

In SBU processing, a single core is circuitized, using conventional board technology and ground rules, by drilling and plating as though it were a finished composite (Figure 2). It is then used as a base for build-up of denser circuit layers. Drilled holes are filled and planarized using a mechanical grinding process. A layer of dielectric is deposited symmetrically on the top and bottom sides of the core, maintaining mechanical flatness through balanced stresses, either as liquid or as dry film, and cured. Vias are formed by photoprocessing or laser drilling through the deposited dielectric layer to the underlying circuitry. Circuitry is formed in a semi-additive process. The entire polymer surface, both dielectric and via hole walls, is seeded with a conductive material in a process derived from conventional circuit board processing, and a thin layer of electroless copper is deposited. A photoresist layer is deposited, patterned, and etched to expose the electroless copper layer and via holes in the desired circuit pattern. Electroplated copper is added to the desired finished metal thickness. Finally, the sacrificial dielectric layer is entirely removed, initial seed metal is stripped, and the SBU substrate is ready for the addition of the next layer pair. This process is repeated for the desired number of layers.

Factors limiting use

At its inception, SBU laminate substrate technology was plagued by a number of limiting factors, the most significant of which was the geometrically compound impact of per-layer yield losses. Although the per-layer

Core	Build-up layers	Finish
Drill	Roughen surface	Roughen surface
De-smear (etch)	Laminate dielectric	Apply solder resist
Plate copper panel	Drill laser vias	Expose resist
Form black copper oxide	Roughen surface	Develop resist
Plug vias	Electroless copper plating	Cure resist
Grind plugs	Apply photoresist	Apply metal finish
Pattern	Apply electrolytic copper	
	Strip resist	
	Etch	

Figure 2

Sequential build-up (SBU) process flow.

wiring capability of SBU is large, other factors such as the need for *x-y*-directed signal wires or power-plane referencing drive the use of additional build-up layers. Whereas substrate yield for a single layer, say 80%, may be reasonable, the compound yield of a three-layer sequential build-up at the rate of $(0.8)^3 = 0.512$ becomes problematic. For this reason, initial designs were simple. Per-layer yields at first were significantly below 80%. To date, there remains a practical limit of six layer pairs (six individual layers on top and six on the bottom) imposed by this compound yield detractor and accumulating yield losses due to the flatness variation resulting from the build-up of unreinforced dielectric layers.

Volume production

Several suppliers were cross-licensed to produce substrates based on SBU technology. Major investments were made in the establishment of high-volume facilities, and the excess production capacity thus created established the beginnings of a commodity market. This available capacity presented an opportunity to IBM and the rest of the industry to migrate applications to SBU.

Rapid breakthroughs in laser drilling technology for micro-vias, initially by Hitachi Industrial in 2000, resulted in significant reduction in cost per micro-via [11]. Practical SBU laminates based on vias formed by laser rather than photoprocessing were first developed by Shinko Electronics Industries [12]. Laminar processing led to a new set of dielectric materials with improved material properties such as low thermal expansion coefficients and high strength at solder reflow temperature (approximately 220°C). Via formation by laser enabled the use of polymer dielectrics containing high loadings of silica particle filler to provide enhanced thermal, mechanical, and electrical properties. This processing allowed assembled semiconductors to be burned-in at higher temperatures without degrading the substrate, in turn enabling a shorter process to eliminate early failures.

These enhanced mechanical properties permitted the use of a surface-mount-pinned interconnect that made possible the first direct socket application of this substrate technology [13].

Additional suppliers were licensed for laser-drilled SBU technology. This approach was independently discovered by other developers, notably Kyocera Corporation and IBM. By 2002, there were more than five companies in mass production, and a commodity market was maturing. The technology industry crash at that time created a buyer's market, and substrate prices dropped precipitously. It became advantageous for an increased range of customers using high-complexity ceramic substrates to investigate the cost-reduction opportunities to be derived from converting planned designs to this low-cost technology.

The flip-chip organic substrate industry evolved around narrow-bandwidth, low-frequency microprocessor designs with modest signal counts. For these applications, a C4 footprint known as a perimeter array was appropriate. Signal C4s require a dedicated substrate trace (an etched copper conductor of narrow width). In a perimeter array, the C4s are densely packed in the outer rows of the array, while power and ground connections occupy inner rows or are interspersed between signal rows. The center of the array is depopulated, which maximizes the number of signals that can escape (i.e., fan out) from the chip C4 cage into a minimum number of dedicated signal planes, usually one. The disadvantage of perimeter array designs is loss of signal quality, which results from increased distance and constriction resistance between a signal and its return path, increasing noise sensitivity.

Implementing IBM designs

The IBM logic C4 footprint designs implemented in multilayer ceramic utilize large numbers of signal planes. Signal traces are wide compared to the chip C4 bump pitch. A single signal can escape the chip footprint only horizontally between a pair of vertical power and ground via connections per layer, and the incremental addition of signal planes results in a proportional price increase. The principles of electrical design can easily be followed by surrounding each signal C4 with a cage of power and ground C4s, which provides the best possible electrical performance. When traditional ceramic-oriented designs are adapted to applications in SBU laminate substrates, the limitations imposed by a single signal plane become obvious. A C4 footprint known as full array was used to surround each signal C4 with a cage of reference voltages. In a full array footprint, signal connections are evenly distributed across the array at minimum density. The array center may be fully populated with signals. In wiring a full array with a limited number of wiring

planes, one finds channels for the fan-out of signal wiring blocked by vias for the connection of power and signals to planes above and below. Fan-out of a similar number of signals utilizes significantly more wiring layers than in a perimeter array.

To adapt its multilayer ceramic-based logic designs to SBU laminate substrates, IBM required layer counts well beyond the set of configurations in which yields were optimized and profitability established.

A competitive market situation resulting from surplus capacity put in place during the technology bubble economy and left idle by the downturn caused a significant change in SBU laminate substrate market conditions. In 2002 quotes showed significant price reductions (of the order of 50%). Given the need to align with customer strategies, and coupled with anticipated performance advantages, this price reduction created an environment favorable for the migration of IBM ASIC designs to SBU.

Alternative technologies

Laminate technologies other than SBU have been proposed to accommodate increasing demands for pin count, power dissipation, and operating frequency. Technologies involving parallel processing, similar to conventional boards but using blind and buried microvias in layers, have been explored, notably Any Layer Inner Via Hole (ALIVH) by Panasonic (Matsushita Electric Industrial, Inc., Kadoma City, Japan) [14]. ALIVH has been in mass production for up to ten-layer circuit boards for almost a decade; alternatives are being explored to adapt it to high-density flip-chip substrate applications. An example of a parallel-process build-up structure obtained from Kyocera Corporation is shown in **Figure 3**.

In 2002, Intel proposed an experimental bumpless build-up layer (BBUL) technology that is characterized by the absence of a conventional core and a direct extension of the outermost metallization layers of the die into the overall thin substrate by eliminating bumps (C4s) on the die [15]. This technology provides the advantages of small electrical loop inductance for power delivery and minimizes discontinuities for high-speed signaling. It also allows for reduced thermomechanical stresses on die materials, high lead count, and ready integration of multiple components such as decoupling capacitors. It has yet to be commercialized.

Description of current technology

An SBU laminate substrate is composed of three distinct technology elements: the surface finish for soldering and adhesion, build-up layers that contain most of the wiring, and a core, which provides mechanical strength (Figure 1). Each element must satisfy widely differing

demands and can be characterized by unique parameters (see Tables 1 and 2, shown later). A numerical naming convention has been adapted to describe the structure of build-up laminate substrates. In this convention, the substrate is viewed in cross section, and the build-up and core metal layers are counted from the top or from the bottom. The practical range of applications varies from 1/2/1, a single build-up-layer pair over a two-metal-layer core, to 6/6/6, six build-up-layer pairs over a six-metal-layer core subcomposite. Most volume applications are for 2/2/2 or 3/2/3 structures.

Surface finish

Surface finish is both an epoxy dielectric, which functions as a solder mask and metal migration barrier, and solder-bumped metal pads on a pattern matching that of the chip C4 connections. Pad pitches in current production are as low as 200 μm , and solderable surface sizes are optimized to the C4 connection, typically somewhat less than the C4 height, either 100 μm or 125 μm . Solder pad areas may be either solder-mask-defined (SMD), in which the solderable area of a larger metal pad is defined by the absence of a solder mask covering, or non-solder-mask-defined, in which the solderable surface is defined by the area of the pad itself. Most current production applications are SMD. It is thought that, to achieve pad pitches in the 180- μm range, SMD will be the only alternative. In fabricating the finish layers, copper pads are formed by the typical semi-additive plating process [16]. The solder resist is either rolled on as a dry film or curtain-coated as a liquid. Solder pad openings are formed through photoimaging. An anti-oxidation finish is applied to the copper, which is typically a two-step process using 1) electroless nickel and 2) immersion gold, although an organic surface protection chemical may be used. Solder/flux paste is screened onto the pads in a process similar to that used in surface-mount assembly. This solder paste is then reflowed to form an adhered metal deposit, and finally coined through a stamping process to flatten the surface for placement of the chip.

Build-up layers

Build-up layers are characterized by copper trace dimensions such as width, thickness, and spacing. Almost all signal wiring in SBU occurs in build-up layers. Also important are dielectric characteristics, thickness, electrical properties such as dielectric constant and loss tangent, and thermal expansion characteristics. Dielectrics are silica-filled epoxies applied as dry films. Via characteristics are also critical. Drilling is through a single dielectric layer at a time; all vias in the build-up layers are blind and buried. Laser-formed vias are tapered, having different lower and upper diameters. The via land diameter is critical because it imposes boundaries

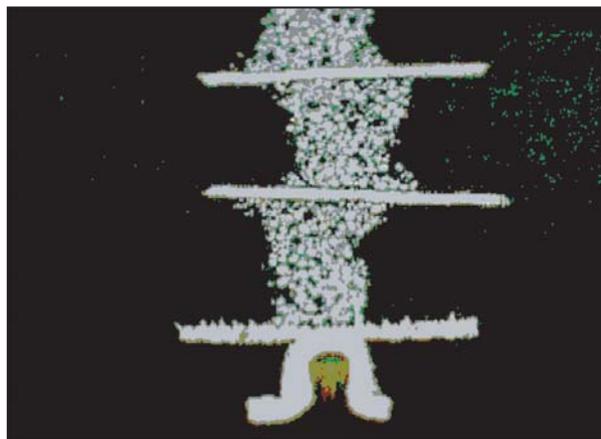


Figure 3

Example of metal-filled epoxy vias connecting copper planes in a parallel-process-fabricated high-density interconnect structure.

on wiring capability and is an indication of the accuracy of the fabricator registration. Via size and pitch is an area of constant focus. Vias have traditionally been stair-stepped, ascending through the build-up layers through the use of lozenge-shaped copper islands bearing the base of one via at one end and the drilled top of a via at the other. In stair-stepped vias, the plating thickness is approximately uniform as deposited both on sidewalls and base. Because stair-stepped vias impose a large blockage to wiring, the focus has been on development of stacked-via approaches. In stacked vias, the via is ideally filled with copper during plating, so that the metal surface at the via top is flat. The via for the next layer is placed with its base directly on the center of the one below. This structure offers significant wiring enhancement as well as enhanced heat transfer, but it is more rigid than the structure formed with stair-stepped vias and readily transfers stresses imposed by differential expansion. Thus, stacked vias are limited in the numbers of consecutive layers to which they can be applied, and are the subject of considerable development focus. At IBM, finite-element modeling is being used to identify necessary materials properties. Once appropriate materials are available, prototypes are fabricated and evaluated through thermal cycling. The ability to withstand 1,000 cycles between -55°C and 125°C is the criterion for success.

Core

The core is composed of glass-fiber-reinforced epoxy, jacketed in subtractively circuitized copper sheets. Most commonly, cores consist of a single dielectric layer, although multilayer cores may be used, formed by conventional laminated printed circuit board processing

Table 1 Typical materials and specifications for SBU laminate.

<i>Package design specification</i>	<i>Material or value</i>
Conductor material	Copper
Conductor resistivity	2 $\mu\Omega$ -cm
Build-up material	Silica-filled epoxy
Build-up dielectric constant	3.1
Core material	Bismaleimide triazine
Core dielectric constant at 1 GHz	4.2
Core line width – minimum	75 μm
Core line width – tolerance	5 μm
Core line thickness – nominal	400 μm
Core line thickness – tolerance	60 μm
Build-up line width – minimum	20 μm
Build-up line width – tolerance	5 μm
Build-up line thickness – nominal	15 μm
Build-up line thickness	5 μm
Build-up dielectric thickness	30 μm
Build-up dielectric thickness	8 μm
Minimum spacing (build-up)	20 μm
Minimum spacing (core)	75 μm

techniques. Where multilayer cores are used, added metal layers are typically used for power distribution rather than signal routing. Copper trace and space dimensions in the core have been significantly coarser than those in the build-up layers. Core vias are formed by mechanical drilling. In the past, the minimum drill size has been 200 μm , and minimum core via pitch with no circuit trace in the 400- μm range. To accommodate a circuit trace for maintenance of a power grid, minimum core via pitches in the 550- μm range have been the norm. Core vias are typically filled with particle-filled epoxy and then mechanically machined flat before build-up layer processing. This establishes the degree of flatness necessary for success in subsequent processes. Where micro-vias stacked over core vias are used, the core vias are plated shut. With C4 pitches at 200 μm easily accommodated by the wiring capability in the build-up

layers, core via pitch has imposed a severe restriction on wiring capability. Thus, in essence, all wiring must be done in the top build-up layers of the substrate. The bottom build-up layers, although processed at the same cost and with the same techniques as the top, have been essentially vestigial, a processing artifact, their sole function being to connect the core vias to the module BGA pads. However, recent developments in mechanical drilling technology have produced a significant advance in core capability. Doubling of drill angular velocity permits drill sizes to be reduced to 100 μm for core thicknesses up to 400 μm with good resultant hole wall quality. This enables core via pitches adequate to fan out half of the signals through the core in the bottom build-up layers. It essentially doubles the wiring capability of a substrate at an incremental cost of more aggressive core ground rules. These fine-pitch cores are required only for applications of the highest complexity.

Typical materials, dimensions, and tolerances for an SBU laminate application are defined in **Table 1**. These values were used for the electrical analysis of the next section. The evolution of characteristic dimensions over time for build-up layers and cores is shown in **Table 2**. An overview of SBU technologies has been published [17]. In addition, for the following illustrative ASIC design case, the package pins can be classified in six groups: power, ground, serial embedded differential, source-synchronous single-ended, source-synchronous differential (which may be edge- or broadside-coupled, parallel or offset), and common clock single-ended (**Figure 4**).

Design challenges for ASIC applications

Emerging ASIC applications possess complex requirements for silicon, package, and end-user system design. Early engagement and interaction between the chip and package designers, and integration of their tools, are essential to addressing these requirements, improving design turnaround time, and making effective cost/performance tradeoffs.

Table 2 Major parameter roadmap for build-up and core layers.

	2003	2004	2005	2006	Comment
Build-up					
Line/space	25/25	20/20	17/18	15/15	15 μm Cu
Via drill/pad	60/95	50/85	40/75	40/75	w/stack
Min via pitch	180	150	120	120	
Core					
Line/space	75/75	65/65	50/50	50/50	15–30 μm Cu
Via drill/pad	150/285	120/200	100/185		0.4-mm core
Min via pitch	325/475	275/450	225/335		0/1 trace between vias

Chips can no longer be designed without knowledge of packaging capabilities. For example, package power grid design is becoming more demanding with low core voltages, which in turn places demands on the package to minimize the IR drops and effectively deliver power. Additional examples of ASIC device trends include signal isolation as well as C4 pitch reduction as a way to increase the usable wafer area, thereby reducing cost. Complexity can be added to the laminates by using more aggressive ground rules such as tighter via pitch, stacked vias, reduced line widths and spaces, and reduced core thicknesses. These trends ripple downstream to the module and assembly manufacturers. Advances are required in process and materials in order to handle thinner, larger carriers, and advances in thermal interface materials in order to manage the higher power densities.

General considerations

The key attributes of organic laminate package technologies as they pertain to the electrical performance of the module are the following: highly electrically conductive metallurgy to minimize IR drops and to effectively deliver power to the chip; low-inductance connections to reduce simultaneous switching noise; low-dielectric-constant insulator materials to better match board impedances and to reduce undesirable parasitic capacitances; and advanced thermal interface materials to manage high power densities on the chip and to improve performance. For high-speed applications, consideration must be given to simultaneous switching noise, electrical coupling noise, signal trace resistance, low dielectric loss, and signal trace characteristic impedance. Depending on the off-chip signaling technique used for the various interfaces on the chip, either near-end or far-end crosstalk noise is a concern. For instance, near-end crosstalk noise is important for full-duplex serial links in which the transmit and receive macros are placed on the chip close to each other because of the sensitivity of the received attenuated signal to coupled noise. Similarly, simultaneous bidirectional signaling is very sensitive to near-end noise as well as to impedance mismatches between the substrate and the board [18]. Therefore, this type of interface requires a controlled-impedance environment, preferably close to the impedance of the transmission link, with low crosstalk and low resistance.

Floor planning

Laminate packaging offers several advantages. It also offers challenges in dealing with some attributes that are more restrictive than those of competitive technologies such as ceramic. It forces a tighter relationship between the chip and package design. The first challenge is to limit the number of build-up layers, thereby reducing package complexity, risk, and cost. Package wiring can help to

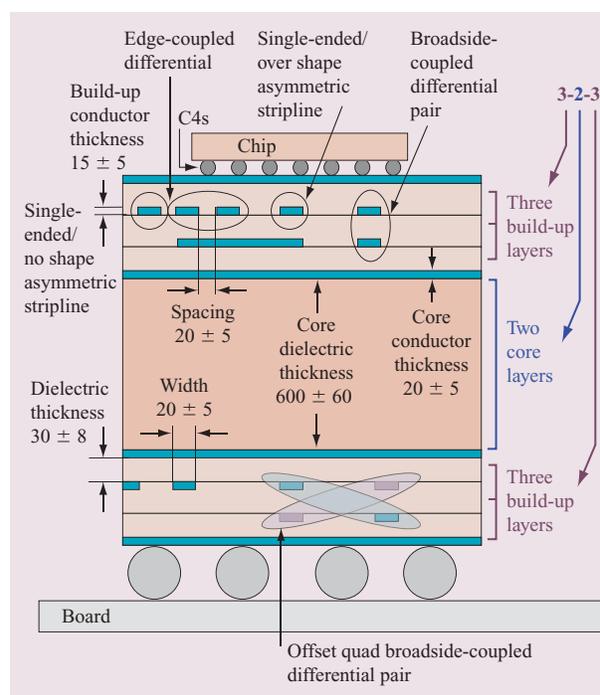


Figure 4

Package dimensions and routing structures for 3-2-3 SBU configuration (all dimensions in μm).

influence the placement of the signal I/O and logic macros, and is influenced by attributes such as required signal/power ratio, core thickness, line width/space, and via density. With an understanding of the package wiring tradeoffs associated with these attributes, this early interaction will result in a balanced solution. There is limited flexibility to add additional build-up layers as the co-design proceeds, if required. Similarly, with the high-performance requirements of the application, jitter is a parameter to be considered. Since the placement of cores such as phase-locked-loop (PLL) circuitry and high-speed serial (HSS) links has a critical dependency on the electrical resistance of chip and package signal traces, the image/package co-design effort must ensure the best possible design in order to minimize the resistance. With the increase in logic cell density of the silicon technology and the higher performance requirement, there is a greater power demand on the package. Compounding that problem is the trend toward lower V_{dd} supply voltage, which allows for less IR drop due to current, so the power grid design is crucial and cannot be sacrificed for an improvement in package wirability.

Modeling and simulation methodology

Simulation methodology for chip/package co-design and optimization involves examination and balancing

of all of the various signaling, timing, noise, and power distribution requirements of all interfaces to reach the performance goals within the cost requirements. Typically, this analysis has two main foci: signal distribution and power distribution. In addition, differential and single-ended high-speed nets, common clocks, source-synchronous signals, and embedded clock timing schemes are common in ASIC designs requiring package design tradeoffs. Driver types include push-pull, open drain, and differential.

With slower speeds and for ease of design, common clock signals require no length matching, allowing for flexibility in routing other length-constrained interfaces such as the source-synchronous and asynchronous interfaces. For source-synchronous interfaces, the clocking signal is sent along with the data; this imposes relative length matching within the group. Matching among groups is not necessary, since each group is clocked independently. Asynchronous serial interfaces require no matching within the bus except for matching within each differential pair. Circuitry at the receiver allows for automatic de-skewing of the data bits for ease of routing.

Impedance control considerations

A critical design parameter for high-speed digital package designs is signal trace impedance to avoid unwanted reflections. Typical digital component applications have requirements for both single-ended and differential signaling interfaces on the same chip. This requirement imposes some restrictions on the trace width and trace spacing, given a fixed dielectric thickness and conductor thickness. For single-ended signals, the trace width can be kept at a minimum, while the trace spacing must be increased in order to satisfy the coupling noise requirements. Typically, this coupled noise specification requires a 3:1 trace spacing-to-reference plane spacing ratio. For these interfaces, the series dc resistance is also a concern and must be kept below a target maximum depending on signal loss allocations. For single-ended signal traces, stripline structures are used to better control impedance and to minimize coupling effects by supplying good current return paths. Differential signaling requires a tightly coupled signal pair (Figure 4) to reduce crosstalk noise and to maintain controlled impedance. For these interfaces, series resistance is not as important, since they have drivers for signal link loss. Near-end crosstalk is important for high-speed serial links, since the transmit (TX) and receive (RX) circuitry are usually located close together on the chip, providing an opportunity to couple noise from the sending signals onto the attenuated receiving signals.

Modeling with a Star-RCXT** three-dimensional field solver (Synopsys, Mountain View, CA) for impedance

indicates that a 50- Ω single-ended transmission line can be accomplished in a symmetric stripline using a 22- μm -wide trace, while an asymmetric stripline would require a 37- μm -wide trace. A 100- Ω differential, edge-coupled symmetric stripline can be constructed using the minimum supported technology width of 20 μm and an in-pair spacing of 60 μm . Routing and noise isolation can be improved by using a tighter spacing (50 μm instead of 60 μm) at the cost of a lower impedance.

With manufacturing tolerances, the design for 50- Ω single-ended impedance traces with 22- μm trace widths can vary considerably. A plot of the impedance variation given a trace width variation of $\pm 5 \mu\text{m}$, dielectric height tolerance of $\pm 8 \mu\text{m}$, and trace thickness variation of $\pm 5 \mu\text{m}$ is shown in **Figure 5(a)**. From the data, a nominal trace width of 22 μm yields a single-ended trace impedance range of 37–66 Ω under worst-case manufacturing conditions.

An optimized stack-up for 100- Ω differential links yields a 20- μm -wide differential pair separated by 50 μm for improved wirability. Geometry tolerance effects are plotted. A permutation of trace width (15, 20, 25 μm), dielectric height above and below the trace (22, 30, 38 μm), and trace thickness (10, 15, 20 μm) yields the distribution shown in **Figure 5(b)**. From the data in **Figure 5(b)**, a nominal trace width of 20 μm yields a differential impedance range of 70–125 Ω under worst-case manufacturing conditions.

Despite the trace impedance variation for both single-ended and differential structures due to manufacturing tolerances, substrates are typically tested to be within a specific impedance range to ensure that the process remains within the required specification.

Noise considerations

Modeling of via coupling can be accomplished using two methods: a quasi-static field solver, which assumes no coupling between the time-varying electric and magnetic fields, thus simplifying Maxwell's equations and reducing the complexity and solution time, and a full-wave solver, which must solve Maxwell's full set. From this, the s- or scattering parameters for a multi-port linear network are extracted for the same geometry. The via pattern and routing extractions are very complex and can be used for post-layout verification. However, they are not suitable for exploratory and design phases because of their size and complexity. Simplified but representative models, such as the one shown in **Figure 6** facilitate the studies through parameterization and speed of execution.

The model signal assignment shown in **Figure 6** is solved using a full-wave field solver (Ansoft HFSS**). The simulated time domain reflectometry (TDR) waveforms corresponding to the coupled noise from

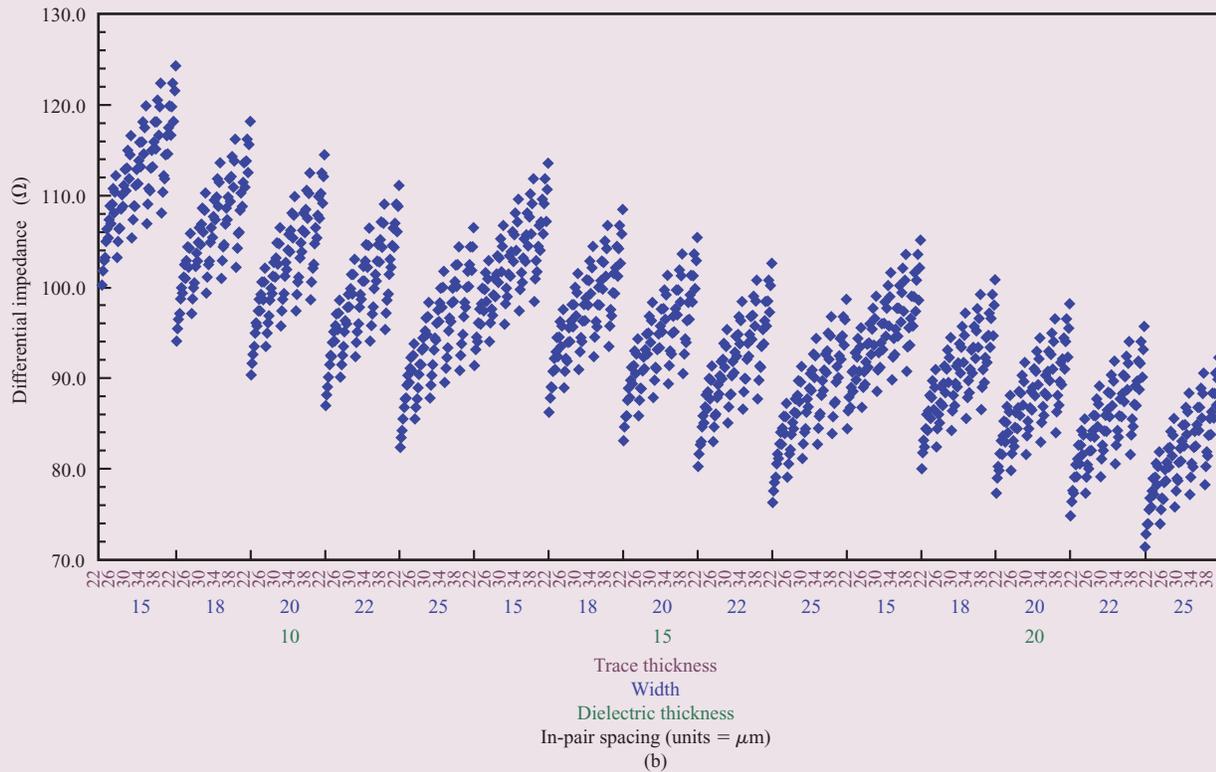
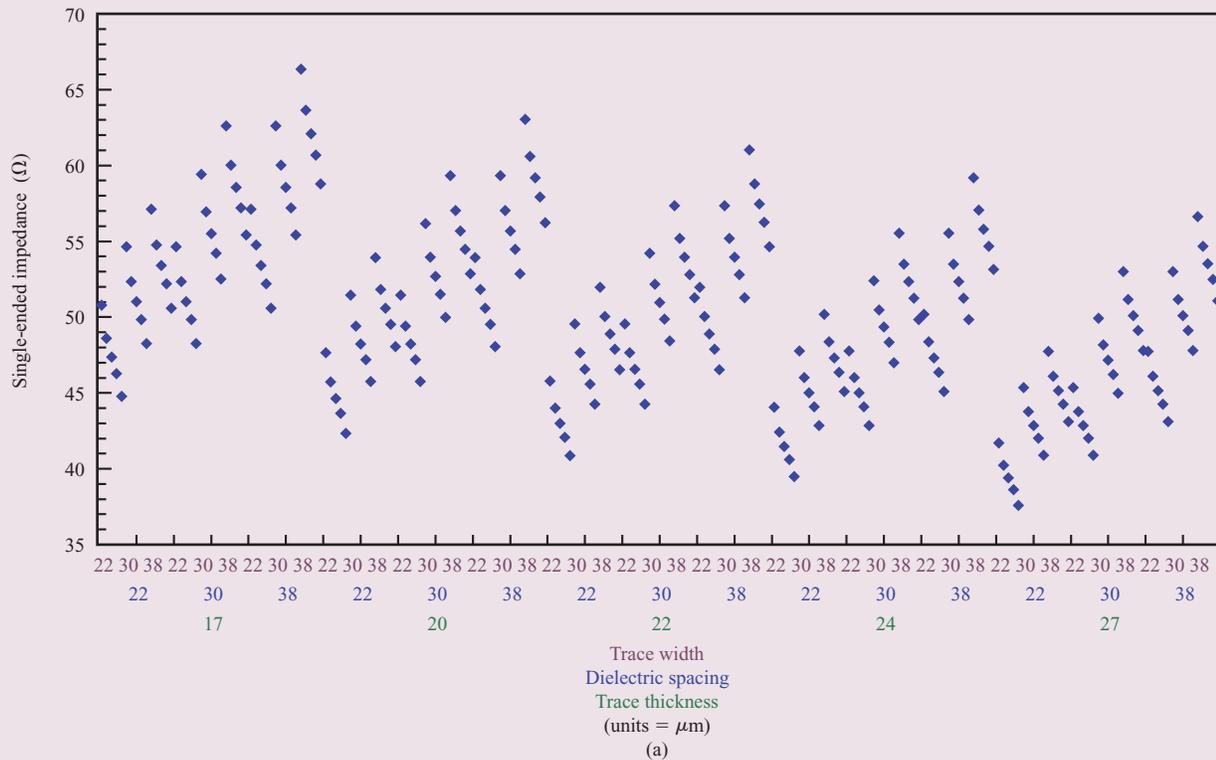


Figure 5

(a) Single-ended and (b) differential impedance variation.

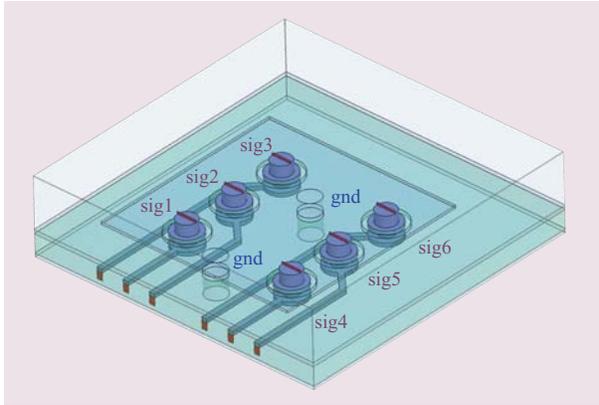


Figure 6

Signal assignment for C4 study.

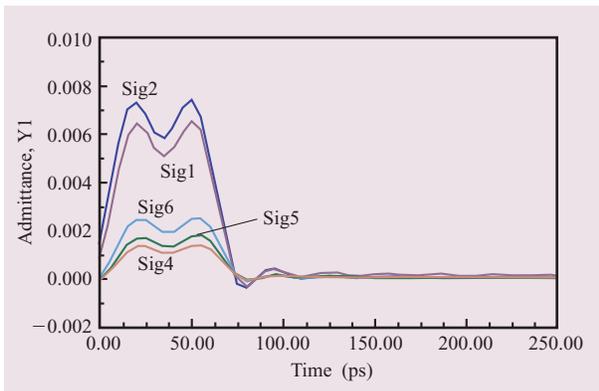


Figure 7

Far-end crosstalk from signal 3 trace to all other traces. Y axis measures coupled noise.

sig3 to sig1, sig2, sig4, sig5, and sig6 are shown in **Figure 7**. The time domain waveform is simulated from the s-parameters extracted with the full-wave tool. In a single-ended case, these values could add together; in the worst case, if sig1 and sig2 were a differential pair, the noise would be common mode and would be rejected by the receiver. Similarly, the noise imposed from a differential pair sig1/sig2 would have minimal far-end noise impact on sig3, since the two noise waveforms would almost cancel each other. Coupling from one row to another is fairly small (less than 0.5%). Full-wave modeling of via coupling indicates that the far-end noise contribution of the non-transverse electromagnetic mode (TEM) propagation in the C4-to-package trace transition is minimal (<0.9% for same-column signals and <0.5% for signals from adjacent columns). This

indicates a good isolation (>40 dB from individual single-ended signals); and the results are even better for differential signals. The model included the top of the C4 to the two stripline routing layers; it did not include the trace on the package, since the propagation mode in this area is predominantly TEM and can be efficiently modeled as a transmission line.

Core via coupling was handled using the same approach as in the C4 region. Using the design rules and actual design implementation geometries, several test cases were modeled using a full-wave approach for maximum accuracy. Solder ball transitions to the board and via coupling are included in the model for the system board. Since card thickness affects via coupling, several card thicknesses must be checked. Typical values range from 70 to 200 mil (1.78–5.08 mm) in thickness.

With minimal far-end noise being generated, the near-end noise dominates the noise in the package. With interfaces with signal rise time in the 60–100-ps range, the near-end noise saturates for traces 5 mm (for 60 ps) to 8 mm (for 100 ps) in length. The maximum trace length observed in a 42.5-mm × 42.5-mm package was 24.2 mm. Near-end crosstalk simulations used a four-aggressor, one-victim trace configuration. Each aggressor trace is driven with a 2-V source step and a 50-Ω source impedance; each aggressor is terminated with 50 Ω, and the victim is terminated at both ends with 50 Ω. The noise induced in the victim trace due to the input step on the four aggressor traces is simulated. To reduce the impact of near-end noise, trace grouping of nets minimizes this effect by routing drivers next to drivers and receivers next to receivers. In cases in which the response trace is not terminated (such as a receiver input), the saturated near-end noise is twice as high, since the near-end noise wave doubles as it reaches the high-impedance end and reflects back into the line. Signals propagating in stripline traces with homogeneous medium have very little far-end crosstalk noise. There are regions where it does not propagate in the TEM mode (e.g., via transitions, which then cause far-end noise). As shown in previous sections, via coupling for this design structure was found to be negligible.

Power distribution considerations

Power plane splits and assignments must be carefully chosen in order to avoid current return path issues. These discontinuities can cause unwanted reflections and coupling and can thus diminish noise margins. Current return path considerations for push/pull drivers are taken into account when designing power plane via locations for current return path continuity. Via distribution for power and current return is very important for high-speed designs. In the design process, an internally developed IBM tool is used as a visualization aid to quantify and

visualize the via distribution from each package layer to the next. The goal is to maintain the uniformity of distribution and number of vias across the entire substrate structure. This number is not less than the number of power and ground C4s on the chip throughout the build-up layers. Due to size and density limitations on the core layer, the number of vias on the core layer must be greater than or equal to the number of power ball grid array (BGA) solder balls.

The frequency analysis of power distribution typically focuses on particular frequency bands depending on the type of analysis and the structure of interest. For on-board decoupling determination, analysis of effective impedance of power may go to a few megahertz. On-package analysis may extend to tens of megahertz because of the lower inductance of typical package capacitors such as interdigitated capacitors (IDCs), low-inductance chip array capacitors (LICAs), and buried capacitance. Since extremely complex models, while possibly having good accuracy, can take a very long time to run, tradeoffs between the model complexity and accuracy must be balanced against the allowable simulation time.

While a full-package model may capture effects such as signal-to-power pin ratio and trace and via coupling, these effects can be studied separately and budgeted for most of the present operating speeds. A simplified model can provide a quick simulation that excludes the various sources of coupling and noise to study signal path effects such as reflections and attenuation.

Design checking

Today's high-performance requirements produce very complicated designs, with thousands of nets, several wiring layers, and numerous electrical criteria that must be defined and checked (impedance, skew requirements, noise coupling, trace resistance, and more). A convenient way to define each of these nets and group parameters is required. System and chip designers, working with the IBM electronic design automation (EDA) group, have developed a series of internal tools called the package-on-the-fly (POF) tool suite. These tools can take chip netlist and driver-type information directly from ChipBench, an IBM-proprietary ASIC floorplanning tool, or they can be used by themselves. They ease the effort of putting package requirements into a format that can be used to automatically design the chip image and package wiring, and develop the circuit-checking routines used to validate the design. Since these tools are integrated, if the package netlist is changed, updates to the package design requirements, chip image rules, and electrical checks are quick, automatic, and synchronized. The electrical checking tools reduce turnaround time by automating the checking process. As more electrical checking is required

for design characteristics such as via density, current return paths, and other wiring details, it will become even more important to have automated tools to validate that the design meets customer requirements.

Design challenges of high-speed, high-power applications

Future microprocessor characteristics

Microprocessor applications represent the major challenge for SBU laminate substrate technologies. There are several key driving factors that have a significant impact on laminate technology. One of the most important of these is the increase in chip size from the 260–350 mm² used today to 260–400 mm² in 2006. Another important factor is microprocessor power, which will increase from 120 W to 200 W in the same time frame [3]. High chip power values will lead to increased current densities that will exceed allowable limits if the design of the laminate is not analyzed and adjusted very carefully. As current density increases, the question whether laminate material will withstand the anticipated higher operating temperatures or break down or degrade is unanswered.

Higher power values together with fragile low-*k* dielectric layers on the chip further accelerate the critical coefficient of thermal expansion (CTE) mismatch problem between the underfill (CTE ~ 50) and chip (CTE ~ 3). New underfill materials used to balance thermal expansion between the chip and the laminate substrate will be needed to resolve this issue. The number of signals will increase from 400 to 800 signals in today's applications to future signal counts exceeding 1,200 signals. These higher signal numbers can only be supported with more build-up layers. Smaller microvias will be required owing to the increased C4 footprint densities, with C4 pitches being reduced from 225 μm to 150 μm. Future microprocessor applications will have more voltage domains. Processor frequencies will increase from 2 GHz to 5 GHz. In conjunction with the higher frequencies, the off-chip bus speed will increase from 0.8 Gb/s to 2.5 Gb/s, requiring impedance tolerances to be lowered from 20% to 10%.

Design considerations

Because of the much higher power dissipation of the microprocessor designs, there is additional complexity beyond ASIC substrate considerations. The design has become a highly complex, multidisciplinary effort.

High-power microprocessors demand high currents through organic laminates, sometimes with current values greater than 150 A. This current must flow without restriction within the substrate, imposing new demands on product functionality and reliability. A power-

management solution implemented anywhere in the substrate starts a design chain reaction that propagates across all layers of the laminate. Power features are key in distributing current, but they also provide return current paths as well as defining tri-plate structures for high-speed signals. Multiple preliminary design optimization cycles through electrical simulation and wiring exercises are unavoidable prior to reaching an acceptable working solution. Signal plane configurations, their number and allocation within the substrate stack-up, and all critical signal paths are affected by power-management arrangements. They must undergo some level of adjustment at every design change.

Laminate substrate applications for high-power ICs in general are bounded by the number of power interconnections that are required to satisfy the current flow requirements from the system board. This is in contrast to conventional thought for chip carriers (i.e., for ASICs), where signal I/Os are the most valued item. In high-power applications with several thousand C4s, at least half of them are power pins, and all of them are required to feed the high current needs of the device. Silicon technology imposes intrinsic limits on the dc-current-carrying capability of chip C4 connections. There is also a dc current limit for each analog BGA connection to the board. Common practice is to allow one 0.6-mm-diameter BGA (with an upper current-carrying limit of 500 mA) for each three C4s (150–200 mA each). The total current values for the various power domains and ground define the total number of required power BGAs. This number added to the I/O BGA determines the size of the module according to internationally standardized fixed BGA pitches. As an example, a single power domain with a 150-A current requirement will require at least 300 power BGAs, with another 300 assigned to the module ground network. In microprocessor applications, the power consumption is more a key factor defining the final module size than simply adding the required number of signal I/Os.

High-current conditions are worse in the substrate areas supplying current to the microprocessor cores. These cores consume most of the chip power, generating areas with very high current density and high temperatures in laminate substrates [18]. These are regions with an absolute requirement to provide low ohmic paths from the board to the chip. By creating several of these low-resistance paths in parallel in the substrate design, dc drops and mid-frequency noise are also minimized.

A few innovative substrate power network design techniques have been developed for handling drops in dc voltage and reducing mid-frequency noise. These start at the very beginning of the project with evaluation and planning of the laminate stack-up with the aim of reducing the number of redesign cycles. Organic

laminates with their copper metallization offer a low-resistance horizontal current path that is not generally available within ceramic packages. The paths represented by the combination of laser vias and plated through-holes that together define vertical circuit transitions are analogous to the wiring paths available in ceramic packages. The horizontal current transfer is achieved through the implementation of full metal layers defined as power planes (PP) within the substrate stack-up. These PP layers, with a copper thickness of 33–35 μm , help in averaging the current distribution across the organic laminate area and are implemented within the rigid multilayer core of SBU substrates such as 4-4-4, a structure consisting of four epoxy-glass-reinforced copper layers jacketed by four unreinforced build-up dielectric copper layers on each side. Each of the main power domains has a PP assigned within the core stack-up.

Other subplanes (with a copper thickness of 15 μm) known as power islands (PIs) are created within the build-up layers to enhance the power distribution network. These are added in all layers including those assigned to signal wiring; good examples are the microprocessor core areas, where very few signals are routed but there is a strong requirement for high current flow. The creation of power islands delivers several benefits for power management but may become a factor limiting the wiring capacity of the signal layers. Power islands have the greatest benefit when these islands are placed as close as possible to the chip; their assignment to a specific power domain is made to optimize the total current handled. The implementation of power islands offsets other undesired effects that would otherwise become disruptive to signal wiring. Without power islands that collect and group power connections, these connections would require independent and redundant vertical structures that are quite rigidly located within the chip shadow area and over their BGA position. A plane or an island permits the relative position of these power vertical connections to be jogged, providing some flexibility in their placement. Once aligned, the wirability of the other layers with their signal-routing channels is preserved. In some designs, PI structures collect and make common almost all of the C4 power connections for the most demanding power domains (ground and core-power). These go directly to a power shape in the top two layers of the laminate, allowing a very effective and more flexible distribution of the power in the inner layers. To further enhance the power-distribution network, these islands and partial planes are then replicated, whenever possible, across the laminate substrate stack-up, establishing connections across all vertical structures belonging to the same power domain. Such parallelization of the resistive vertical circuit structures greatly reduces the final resistive value of the total current path.

The same concept of parallel structures is applied to circuit features such as the laser vias connecting the BGA to the above power plane. In addition to C4s, laser vias and other structures have intrinsically limited current capacity. Specifications must be established limiting current to levels below which structures degrade or cause degradation of organic laminate materials. To meet these specifications and to lower as much as possible the overall resistive value for the current path, the parallelization approach has been applied between all metal circuit features belonging to the same power domain. With this approach, a dc-drop overall value of only 3.95 mV was achieved as the total contribution of all structures (BGA to chip into a 4-4-4 carrier) required to fully power up a microprocessor core using 31 A for the main power.

Another example of the use of PI is the connection of decoupling capacitors, placed close to the silicon with very short connections to large metal planes feeding the chip power grid. The different power planes are placed within the top two layers of the substrate only 30 μm apart from one another. Using these techniques, power interconnections are sufficient to provide return paths for the signal arrays at ratios of either 1:1 (one power BGA to one I/O BGA) or 1:2 (one power BGA to two I/O BGAs).

Power management for high-power applications will benefit from future organic technology enhancements. For example, implementation of better dielectrics, multiple stacked vias, and stacked vias placed on top of plated through-holes will allow the achievement of densities for vertical paths equal to the ones now present in the silicon chip C4 power layouts.

Current microprocessor architectures use single-ended signals. Miniaturization trends, increased currents, and higher and higher signal frequencies may require, in the near future, solutions with better control of noise and signal crosstalk. This improvement will require migration from single-ended signal lines to differential pairs, and presents a complete new set of challenges to the use of organic laminate technology for microprocessors.

Current-induced heating and its implications for design

A consideration that must be addressed when designing organic carrier solutions, especially for high-power applications in which current transfer in the power-distribution network can reach 100–200 A, is the effect of current-induced heating, better known in PCB technology as joule heating. For low-power solutions, this is inconsequential, since current is low and resistance of the individual trace features is small. As power increases, heat generated in these features increases, causing temperature gradients within the carrier that have a direct effect on long-term reliability. Of concern is electromigration of

the conductive material itself, enhanced degradation (primarily cracking) of the dielectric material during power cycling, and electromigration of conductive material through these cracks. An additional concern is the increase in total heat generation within the module, which cannot be allowed to increase the device operating temperature. To compensate for this, enhancement of the module-level thermal solution is necessary to maintain performance as well as reliability targets. While these effects are not currently fully understood, characterization of the local temperature rise as a function of induced current through various features has been measured and confirmed through modeling. Further work is needed to validate the models by measurements on specially designed test vehicles undergoing power cycling. In the interim, design guidelines are based on limits as defined in recent literature. Intel's study of this effect in support of its Pentium 4 processor package designs recommends limiting the carrier temperature build-up to 120°C for areas outside the device shadow [19]. To maintain temperatures below this limit for their family of applications and as defined by their reliability studies, the maximum current densities allowable were 0.33 mA/ μm^2 through any conductive feature.

It is imperative to treat each application as unique. Full thermal modeling of the entire carrier structure, especially for power distribution, is needed. Required via densities in the core and build-up areas are significantly different for lateral and vertical power feed designs.

In developing design guidelines for a particular application dissipating 100 watts, thermal modeling was used to estimate the temperatures resulting from current-induced heating in laminate traces and micro-vias and to establish general relationships between current and temperature in the features. Understanding both the current/temperature relationship and the thermally induced failure mechanisms in the package would allow one to predict reliability. The first thermal model results, shown in **Figure 8**, made it clear that the current/temperature relationship is very strongly influenced by the geometry and boundary conditions. The figure shows the local temperature rise for a 20- μm -wide trace for different depths in the laminate (distance from surface) and different levels of surface cooling. Different heat transfer coefficients (HTCs) were applied on the surface of the laminate, the range of values representing different types of cooling, such as airflow over the laminate or direct conduction through a lid to a heat sink.

Results showing the dependence of temperature and heat transfer coefficient on location in the laminate clearly indicate the need for application-specific modeling, focusing on worst-case scenarios in specific environments.

The processor package design and cooling solution for a particular set-top box application was modeled to find

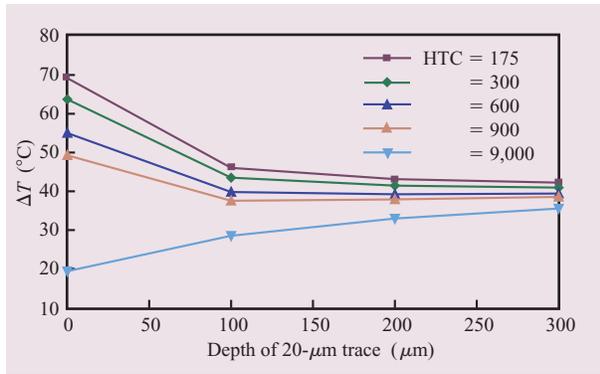


Figure 8

Temperature rise for 20-μm-wide trace vs. depth of trace in laminate for 0.5-A current, for different levels of surface cooling. HTC = heat transfer coefficient (W/m²K).

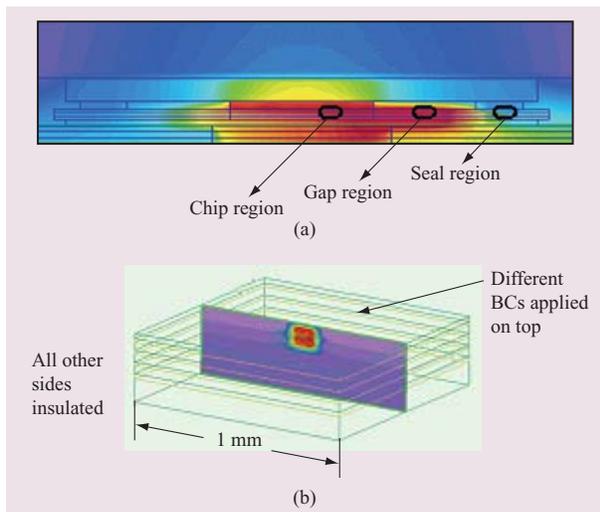


Figure 9

(a) Full package model and (b) micro-via model. The full model was used to extract boundary conditions to apply to the micro-via model. Range of temperature indicated by color (red: 93.4°C–92.6°C; yellow: 92.6°C–91.7°C; green: less than 91.7°C).

the current vs. temperature relationship for a micro-via in the laminate. This is a high-power, high-current application in which current-induced heating and thermally driven failure mechanisms are a concern. The boundary conditions for the micro-via model were obtained from the full model of the package, as shown in **Figure 9**. The full model contained the chip power map with 100 W distributed on the chip surface as well as 3 W uniform dissipation in the build-up layers on the side of the package with the highest current. The base

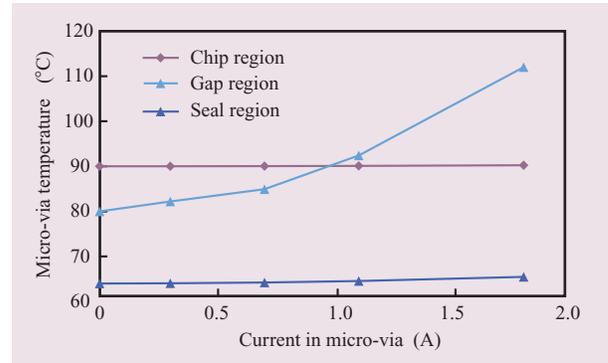


Figure 10

Modeling result for 50-μm micro-via showing peak temperatures as a function of current for three different regions: under chip, between chip and seal (gap region), and under seal.

temperature and heat transfer rate off the laminate surface were extracted from this model for the region between the chip and lid seal (gap region), and the region directly under the lid seal (seal connecting the lid to the laminate carrier). For the region under the chip, a constant surface temperature was assumed, corresponding to the chip maximum temperature. This is a good assumption, since the package cooling solution and thermal monitoring system will be designed to maintain a chip temperature under the specified limit, which in this case is 85°C average. The results are shown in **Figure 10**. The plot shows that the critical region is the gap region, where the heat transfer is very poor and the base temperatures are high because of the chip heating.

The peak temperatures for a micro-via in the critical region can exceed 110°C for currents above 1.5 A. Further evaluation of the current distribution within the upper and lower build-up regions is necessary to verify this assumption. Design rules must be implemented to prevent the occurrence of such high currents. A test vehicle has been designed which is intended to validate characterization as well as to undergo power cycling to establish the reliability relationship with the characterization results.

Thermal challenges in high-power applications

The demand for higher speeds drives the need for smaller devices and higher circuit densities, resulting in increased power density and higher overall power due to leakage. As a result of these conditions, the package substrate cooling requirements become significant. The thermal interfaces from the chip to the metal lid and from the lid to the heat sink, a finned metal structure designed for convective cooling, are typically the highest-resistance

elements in the thermal path. The package designers' challenges are to reduce interface resistances while maintaining mechanical reliability. The challenges are exacerbated by new requirements for lead-free solder joints requiring higher process temperature.

Thermal performance requirements

Although the total power resulting from the increased numbers of circuits per chip and the increased leakage per device are approaching more than 200 W in some applications, the thermal challenge at the package level is driven as much by the high local power densities as it is by the total power. Local power densities are approaching 4 W/mm² and can create significant cooling challenges even when the total chip power is less than 100 W. Processor units on the chip can be significantly hotter than a majority of the chip area and may thus require advanced cooling solutions to keep peak chip temperatures below design targets. Typical peak junction temperature targets ($T_{j,max}$) are between 70°C and 105°C, based on performance requirements. Local ambient temperatures are typically 30°C to 50°C, resulting in temperature budgets typically in the 40–50°C range.

The high-power and low-temperature budgets require very low thermal resistance for the package and external cooling solution. Significant advances in air- and water-cooled heat sinks have been made to meet the total power dissipation challenges, but the focus here will be on the first-level cooling challenges. **Figure 11** shows the typical single-chip package and associated thermal resistance path. The first-level thermal components include R_{jb} , R_{TIM1} , and R_{spread} . Of these, R_{TIM1} (the chip-to-lid interface resistance) is usually the gating thermal resistance. TIM1 unit resistance requirements for high-end applications are typically less than 15–20°C-mm²/W, with some applications capable of less than 10°C-mm²/W. R_{jb} is the junction-to-board resistance, which is usually at least an order of magnitude higher than the resistance of the path through the chip and heat spreader and is usually ignored.

The unit resistance of the interface is the ratio of the bond line (the cured adhesive layer joining two thermal components) to the effective conductivity of the interface material, including the interfacial components between the thermal interface material (TIM) and the two mating surfaces. To lower the resistance, one can reduce the bond line or increase the effective conductivity. Typical interface materials are described in the next section.

Interface material challenges

Figure 12 shows a fundamental problem associated with organic packages. When a low-CTE die (~3 ppm/K) is coupled with underfill to a higher-CTE composite laminate (15–20 ppm/K), the result is a warped package

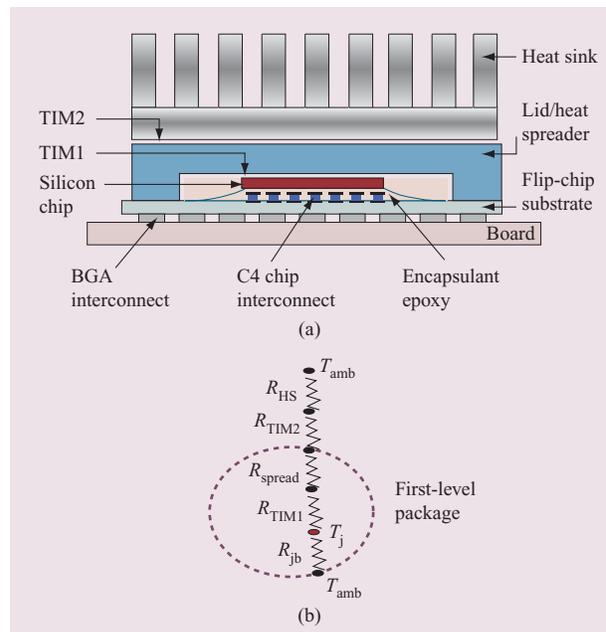


Figure 11

(a) Lidded organic package with typical cooling solution and (b) thermal resistance path. (TIM = thermal interface material.)

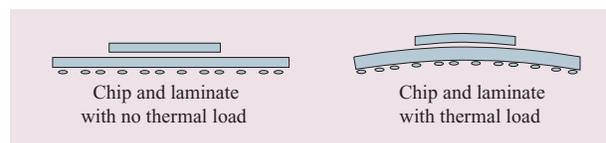


Figure 12

Package/chip warpage caused by CTE mismatch between chip and lid.

and chip. Depending on chip size, the warpage of the chip can exceed 60 μm. When a flat lid is attached to the chip using an adhesive, the outer edges of the adhesive are in significant tension. This not only creates a larger bond line at the edges, but compromises reliability as well. For pastes or greases, cyclic thermal loading can create a pumping of the TIM as the package undergoes significant displacement from the flat to the warped state.

Three categories of thermal interface materials are used in packaging: adhesives, pastes/greases, and phase-change materials. Although similar thermal performance can be achieved with all three categories, the mechanical properties and reliability distinguish one from another.

Adhesives generally consist of high-conductivity filler material such as silver, aluminum, alumina, or copper, embedded in a silicone or epoxy matrix. There are many

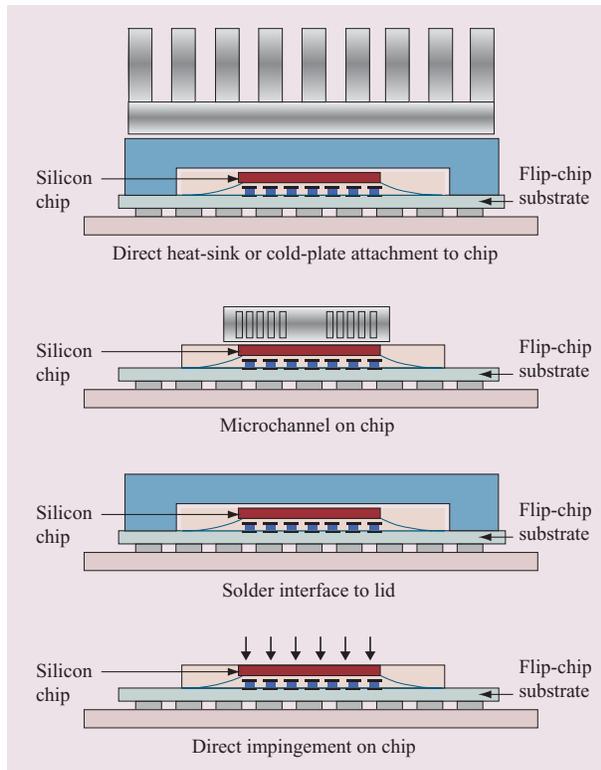


Figure 13

Advanced cooling concepts for organic packages.

commercially available thermally conductive adhesives. Typical effective conductivities in product form factor range from 1 to 4 W/m·K, although higher values have been reported by vendors. Typical bond lines range from 100–50 μm down to 10 μm . A lower-effective-conductivity material can outperform a higher conductivity material if the bond line can be made significantly smaller; however, this requires the mating surfaces to be very flat. The mechanical properties of thermally conductive adhesives can be very different. Silicones offer a low modulus of elasticity (<15 MPa) and have adequate adhesion strength, which allows sufficient decoupling between the die and spreader. Higher-modulus epoxies can be used only with CTE-matched spreaders or at very thick bond lines. The adhesives generally have cure temperatures in the 100–150°C range, and also have to survive high solder reflow temperatures (>250°C for BGA lead-free applications).

Pastes and greases are another category of materials that offer thermal performance similar to that of the adhesives (1 to 4 W/m·K, bond lines 15–100 μm) and generally are reworkable and do not require any curing. One challenge for organic packages with pastes and

greases is that the material can be pumped out owing to the significant package displacements in thermal cycling. Gel materials which have some level of cross-linking but virtually zero modulus are one approach to solving this problem.

Phase-change materials are another category of TIMs with similar thermal performance and reworkability as pastes and greases. Bond lines achieved with these materials are generally made higher by the carrier material, which is typically used to support the conductive phase-change materials; however, excellent thermal performance has been measured with some commercially available phase-change materials (<15°C mm²/W). Reliability through thermal stressing is a significant challenge for phase-change materials.

Next-generation and advanced solutions

To meet the increasing performance and power density challenges, next-generation packages are focusing on highly conductive interfaces, such as solder, and direct attachment of the heat sink to the chip. These solutions have already been explored and made successful by IBM in ceramic packages, where the mechanical displacements are minimal and high process temperatures can be accommodated. Direct heat-sink attachment is also made simpler in ceramic packages because TIM pumping is not as severe. A major obstacle for organic packages is chip cracking resulting from handling damage caused by shipping without a lid or damage caused by the assembly process at the second level. Since the chip is under high tensile stress, a small defect or scratch can create chip cracking and significant yield loss. For organic packages, very-low-modulus solders have been used successfully by others on smaller chip sizes. Solder interfaces on larger chip sizes (>150 mm²) remain a challenge.

Alternatives such as direct liquid impingement onto the chip are also being explored, but these present their own challenges, such as leak-proofing and isolation of the liquid. Another approach is to attach a silicon microchannel to the back of the chip [20]. In this case, the CTE-matched structure can reduce the mechanical stresses at the bond line. Prototypes have already been used on ceramic packages in IBM. **Figure 13** shows some of the advanced cooling concepts for organic packages investigated by IBM.

Socketing of SBU modules

Socketing of SBU modules in high-end servers would facilitate field replaceability, which is desirable on circuit boards with more than one processor module. It also presents additional technical challenges and some added cost per I/O. The most common socket types are land grid arrays (LGAs), hybrid BGA/LGA, and pin grid arrays.

LGAs consist of an array of two-sided compliant contacts held in an insulating carrier. These are available at 1.00-mm pitch and large array sizes from a variety of suppliers. The LGA contact array, interposed between the module and board, requires a constant application of force pressing it between the module and board in order to maintain electrical contact. This force is typically supplied by a mounting hardware assembly consisting of a load-plate/heat-sink combination which overhangs the module laterally. At the bottom of this load plate near the corners are four posts that extend down through clearance holes in the board, through a bottom load plate, and are attached to a spring mechanism that applies the total load. In most LGA types, this force is designed to be between 60 and 80 grams per contact. This magnitude of force applied through the relatively soft organic SBU module is known to exacerbate chip and dielectric cracking, especially with temperature cycle testing. Low-force LGAs in the range of 30 grams per contact are now available and will help address this concern. Also of concern when using LGAs with organic modules is that nonplanarity of the module bottom surface metallurgy (BSM) can lead to uneven distribution of contact force across the LGA. This was shown by Monte Carlo modeling for the case of a large ceramic substrate, where the distribution required compensation in order to be brought closer to target contact force values [21]. This work identified topography and module stiffness as among the most important contributors to this force distribution; both of these are known to be worse in SBU modules than in ceramic substrates. Indeed, 4 to 6 mils of warp is not uncommon. Not only is the magnitude of non-planarity greater than with ceramic substrates, but the shape of a lidded SBU module is more complex as well. These often resemble a volcano rather than the simple concave or convex camber which is most typical of ceramics. Such factors make high compliance in LGAs even more important than in previous applications. **Figure 14** shows the contact surface topography of a lidded and stiffened organic module measuring approximately 47.5 mm \times 47.5 mm. The CTE mismatches of all components and their relative stiffness values combine to produce this shape, with the high points at the chip shadow corners and maximum displacement of approximately 2 mils for this particular module.

Hybrid LGA/BGA is a technology in which each contact is soldered onto the board-side BGA pad and reaches outward with a compliant end that meets the bottom of the module. This is held under constant application of force, similarly to LGAs. Hybrid varieties are being utilized by Intel in the Socket T [22]. One advantage of this over full LGA is that it eliminates the

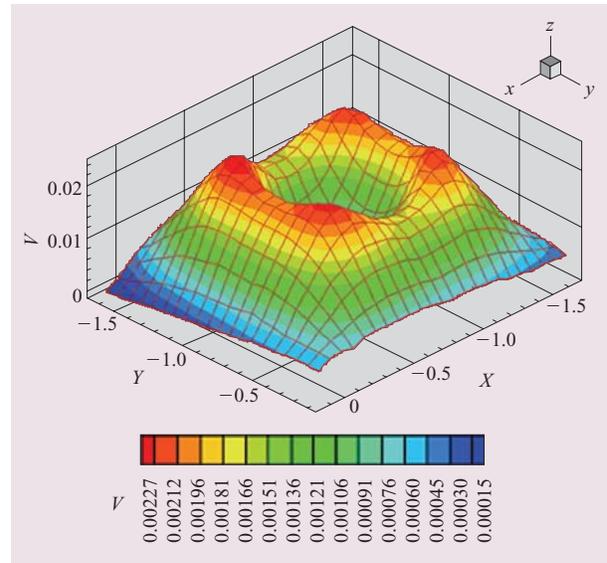


Figure 14

Topography of a 47.5-mm \times 47.5-mm lidded organic module with stiffener. (Courtesy A. Mikhail, IBM Rochester.)

need for costly thick gold metallurgy on the board pads.

Pin grid array at the 1.00-mm pitch is often referred to as micro pin grid array (μ PGA or mPGA). This technology utilizes a zero- or low-insertion-force socket which is soldered to the board. Pins are connected to the bottom of the module. Several companies are developing this technology, which holds the promise of being a reliable connector type. There are currently two main concerns with μ PGA. To date, it is not widely available in large array sizes, though this is a temporary concern. Also, the handling of fragile pinned modules requires great care. Pins can easily be bent during installation; once in place, however, they are likely to be robust. This is more of a concern with low-end modules than with the high-end product space, because such modules are more likely to be maintained by consumers rather than by trained service engineers.

Manufacturing, quality, and reliability challenges and considerations

In this rapidly evolving technology, challenges are compounded with each evolutionary generation. Conventional 800- μ m-thick cores, with via pitches of 500 μ m, are relatively stiff and retain planarity through processing. Cores 400 μ m thick are much more readily affected by asymmetries in circuitry and by processing stresses. Planarity in fine-core substrates becomes a critical manufacturing parameter. Whereas resistance to

conductive anodic filament (CAF) growth in temperature humidity bias stress is well established for hole-to-hole spacings of 200 μm [23], in fine-pitch cores with spacings of 100 μm , resistance to CAF is more difficult to achieve, and the appropriate choice of materials and processes is less well understood.

Interfacial issues between core and build-up layers arise in denser applications. The use of stacked micro-vias over filled-core vias has proven problematic owing to pumping action from expansion of the via fill material. Micro-vias perched over drilled via hole walls show much better reliability performance. When forming stacked vias in the build-up layers, it is virtually impossible to plate a micro-via full so that its surface is entirely flat. A dimple remains in the via top surface. Because of the tapered nature of the laser-drilled hole through the next-layer dielectric, the area of the joint between two stacked vias, and thus the resultant reliability of the stacked structure, is highly dependent on the depth of the dimple. Because of their rigidity, achieving thermal cycling reliability for multilayer via stacks is challenging. The current challenge is to achieve a reliable three-micro-via stack.

For organically based structures, as physical properties of the material deteriorate above the glass transition temperature (T_g), usually in the range of 120–180°C for epoxy, solder process temperature is always an issue. Moisture absorbed from atmospheric humidity during processing and storage vaporizes during solder processing and exerts pressure on interfaces [24], particularly at copper planes where the interface is continuous over a large area. To prevent interfacial separation, planes are perforated to permit moisture to escape. Increasing filler in the dielectric layers to achieve the thermal expansion properties is needed for cycling reliability. This may reduce the adhesive strength of the layers, since filler particles occupy a larger portion of the dielectric-to-copper interface. The migration to lead-free applications, with its attendant significantly higher process temperature, will add additional challenges in this regard.

The interface between substrate and silicon chip, the C4 cage, remains problematic. Differential expansion differences between chip and substrate are large, typically 3 ppm/°C vs. 17 ppm/°C. Increases in die size result in linearly increasing strains. Necessary perfect adhesion of underfill to solder resist after C4 solder processing and through BGA assembly is difficult. Reducing C4 pitches and bump sizes reduces the thickness of the underfill stress buffer and places new requirements on the viscous properties of underfill for gap filling. The migration to fragile low-dielectric-constant materials in the chip wiring layers adds a new demand for stress reduction to the system. Here again the higher temperatures and new material sets added by

lead-free processing are a burden. Substrate barrier and wetting metal layers may have to be modified to accommodate lead-free alloys.

Conclusions

The development trend toward rapidly evolving and increasingly complex laminate packages has been discussed. The capability of SBU technology to handle high full-array signal counts and high chip power levels has been presented. The implementation of high-performance, cost-effective designs for microprocessors and controllers in this technology requires a holistic approach that comprises the total interconnect structure including the silicon, the package, and the system. By properly exploiting the attributes of these regimes, optimal performance and cost can be realized. The importance of properly designing the substrate for high-speed signaling, including identification of key parameters and design tradeoffs, must be emphasized; otherwise the performance and cost of the package are jeopardized.

The key electrical design parameters for high-speed applications are simultaneous switching noise, electrically coupled noise, signal trace resistance, and signal transmission impedance. Laminate packages provide lower dc resistance than ceramic packages. In addition, tightly controlled impedance values are easier to achieve with laminate packages than with ceramic packages, given the dielectric material characteristics and the achievable trace width and trace spacing.

For laminated organic substrate technology and components requiring high-power/high-current applications, a published work has identified a design limit related to joule heating in the substrate circuitry [19]. Future microprocessor designs will place increasing thermal challenges on the technology and will require further invention and development. In addition, further work is needed to determine the power and thermal envelopes required for emerging applications.

It is critical that the processes for the design of the chip top-level metal and the package–chip footprint be integrated to enable designs which can manage the extreme level of complexity associated with emerging silicon technologies. Decisions made or not made early in the product design affect cost, quality, and performance, as well as schedule and time to market. In summary, the rapid evolution of semiconductor applications will continue to place increasing demands on this emerging packaging technology.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Ciba Specialty Chemicals, Intel Corporation, Synopsys, Inc., or Ansoft, Inc.

References

1. International Technology Roadmap for Semiconductors (ITRS)—2003 Edition: Assembly and Packaging, Sematech, Inc., Austin, TX.
2. D. Edenfeld, A. B. Kahng, M. Rodgers, and Y. Zorian, "2003 Technology Roadmap for Semiconductors," *Computer* **37**, No. 1, 47–56 (2004).
3. R. Shukla, V. Murali, and A. Bhansali, "Flip Chip CPU Package Technology at Intel: A Technology and Manufacturing Overview," *Proceedings of the Electronic Components and Technology Conference*, IEEE, May 1999, pp. 945–949.
4. G. Harmon, *Wirebonding in Microelectronics, Materials, Process, Reliability, and Yield*, Second Edition, McGraw-Hill Book Co., Inc., New York, 1997.
5. R. Tummala, E. Rymaszewski, and A. Klopfenstein, Eds., *Microelectronics Packaging Handbook*, Van Nostrand Reinhold, New York, 1989.
6. J. Joly and D. Lambert, "Solutions to the Needs for Higher Interconnect Density on PWB's: Build Up Technology and Micro Vias," *Proc. SPIE* **3830**, 204–209 (1999).
7. Y. Tsukada, S. Tsuchida, and Y. Mashimoto, "Surface Laminate Circuit Packaging," *Proceedings of the Electronic Components Conference*, IEEE, 1992, pp. 22–27.
8. J. R. Millar, "Interpenetrating Polymer Networks. Styrene-Divinylbenzene Co-Polymers with Two and Three Interpenetrating Networks, and Their Sulphonates," *J. Chem. Soc.*, pp. 1311–1317 (1960).
9. M. Nakamura, M. Kato, M. Asai, and H. Takenaka, "High Reliability, High Density Build Up Printed Circuit Board for MCM-L," *International Conference and Exhibition on Multichip Modules (ICEMCM)*, Society of Photo Optical Instrumentation Engineers, Bellingham, WA, 1995, pp. 36–41.
10. V. K. Nagesh, R. Peddada, R. Ramalingam, B. Sur, and A. Tai, "Challenges of Flip Chip on Organic Substrate Assembly Technology," *Proceedings of the Electronic Components and Technology Conference*, IEEE, May 1999, pp. 975–978.
11. R. Schaeffer, "Trends in Laser Microvia Drilling," *Printed Circuit Design & Manuf.* **21**, No. 5, 34–37 (May 2004).
12. S. Wakabayashi and M. Shimizu, "Technical Trends of LSI Packaging: Recent Advances in CSPs and High Density Substrates," *Digest of Technical Papers, Symposium on VLSI Circuits*, IEEE, June 17–19, 1999, pp. 69–72.
13. H. Yeoh, M. Lii, B. Sankman, and H. Azimi, "Flip Chip Pin Grid Array (FC-PGA) Packaging Technology," *Proceedings of the Third Electronics Packaging Technology Conference*, December 2000, pp. 33–40.
14. D. Andoh, Y. Tomita, T. Nakamura, and F. Echigo, "The Progress of the ALIVH Substrate," *Proceedings of the 52nd Electronic Components and Technology Conference*, IEEE, May 2002, pp. 1419–1424.
15. H. Braunisch, S. N. Towle, R. D. Emery, C. Hu, and G. J. Vandentop, "Electrical Performance of Bumpless Build-up Layer Packaging," *Proceedings of the Electronic Components and Technology Conference*, IEEE, May 2002, pp. 353–358.
16. C. Yeh, K. Lan, W. Dow, J. Hsu, and C. Lee, "Micro Via Filling Plating Technology for IC Substrate Applications," *Circuit World* **30**, No. 3, 26–32 (2004).
17. I. Ho, "An Overview of SBU (Sequential Build-up): Microvia Technologies," *Proceedings of the International Symposium on Advanced Packaging Materials*, IEEE, 1997, pp. 19–23.
18. D. N. deAraujo, M. Cases, N. Pham, and B. Rubin, "Substrate Design Optimization for High Speed Links," *Proceedings of the Electronic Components and Technology Conference*, IEEE, May 2003, pp. 1419–1424.
19. A. Hasan, A. Sathe, D. Wood, and R. Viswanath, "Pentium 4 Processor Package Design and Development," *Proceedings of the Electronic Components and Technology Conference*, IEEE, May 2003, pp. 1431–1439.
20. E. G. Colgan, B. Furman, M. Gaynes, W. Graham, N. LaBianca, J. H. Magerlein, R. J. Polastre, M. B. Rothwell, R. J. Bezama, R. Choudhary, K. Marston, H. Toy, J. Wakil, J. Zitz, and R. Schmidt, "A Practical Implementation of Silicon Microchannel Coolers for High Power Chips," *Proceedings of the 21st Semiconductor Thermal Measurement and Management Symposium*, IEEE, March 2005, pp. 1–7.
21. J. C. Parrilla, F. E. Bosco, J. S. Corbin, J. J. Loparco, P. Singh, and J. G. Torok, "Packaging the IBM eServer z990 Central Electronic Complex," *IBM J. Res. & Dev.* **48**, No. 3/4, 395–407 (2004).
22. Intel Corporation, 2200 Mission College Blvd., Santa Clara, CA 95052.
23. C. Navarro, "Development of a Standard Test Method for Evaluating Conductive Anodic Filament (CAF) Growth Failure in PCBs," *Circuit World* **28**, No. 2, 14–18 (2001).
24. S. L. Buchwalter, P. J. Brofman, C. Feger, M. A. Gaynes, K.-W. Lee, L. J. Matienzo, and D. L. Questad, "Effects of Mechanical Stress and Moisture on Packaging Interfaces," *IBM J. Res. & Dev.* **49**, No. 4/5, 663–675 (2005, this issue).

Received October 19, 2004; accepted for publication July 6, 2005; Internet publication September 9, 2005

Edmund D. Blackshear *IBM Integrated Supply Chain, 2070 Route 52, Hopewell Junction, New York 12533 (ed.blackshear@us.ibm.com)*. Mr. Blackshear is a Senior Technical Staff Member in the Procurement Engineering organization. He received a B.S. degree in mechanical engineering from the University of Minnesota in 1982, joining IBM that same year, and an M.S. degree in electrical engineering from Syracuse University in 1988. Mr. Blackshear is currently the team leader for electronic packaging within procurement. He is a member of JISSO, an international standards body focusing on packaging issues, and a past JEDEC committee chair. He received the JEDEC Chairman's Award for technical achievement in 2001. Mr. Blackshear has nine issued patents and three pending.

Moises Cases *IBM Systems and Technology Group, 11400 Burnet Road, Austin, Texas 78758 (cases@us.ibm.com)*. Mr. Cases is a Distinguished Engineer in the xSeries* hardware development organization at the IBM Austin facility. He received a B.S. degree in electrical engineering from City College of New York in 1969, an M.S. degree in electrical engineering from New York University in 1973, and an M.S. degree in computer engineering from Syracuse University in 1979. Mr. Cases joined IBM in 1969 and has worked on numerous areas related to the design and development of digital computer circuits and systems. He is currently the team leader responsible for the xSeries high-end and BladeCenter* system electrical design and performance. He also chairs the Electromechanical Working Group (EWG) for the Infiniband Trade Association (IBTA). Mr. Cases has 16 U.S. patents issued and three pending. He has 59 external refereed technical publications on IEEE transactions and conferences. He is a Senior Member of the IEEE and, since 2002, an Associate Editor for the *IEEE Transactions on Advanced Packaging*; he has received three IEEE awards for his outstanding contributions. Mr. Cases has received four IBM Outstanding Innovation Awards for his work in various fields such as microprocessor chip design, circuit design, and system electrical design. He is an IEEE Senior Member and a member of the Tau Beta Pi and Eta Kappa Nu honor societies.

Erich Klink *IBM Systems and Technology Group, Schoenaicherstrasse 220, 71032 Boeblingen, Germany (EKLINK@de.ibm.com)*. Mr. Klink is a Senior Technical Staff Member in the IBM Systems and Technology Group, where for the last 15 years he has worked on the design and electrical analysis of first- and second-level packages for ipz servers. He is currently responsible for the packaging technology of future systems. The areas of his expertise are packaging technology and design. Mr. Klink has written many papers on these subjects, and he holds several patents. He received his diploma in electrical engineering from the University of Stuttgart, Germany, and he is an IEEE member. Mr. Klink is currently mentoring with the University of Torino, Italy, for research applicable to industry needs.

Stephen R. Engle *IBM Systems and Technology Group, 1701 North Street, Endicott, New York 13760 (englesr@us.ibm.com)*. Mr. Engle is a Senior Engineer in the IBM Systems and Technology Group. He joined IBM in 1982 as a graduate of Clarkson University and has worked in numerous areas in the field of packaging development, both as an engineer and as an engineering manager. He is currently responsible for organic chip carrier applications in support of IBM server products. Mr. Engle holds five patents and has numerous publications in his field of expertise.

Ronald S. Malfatt *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (malfatt@us.ibm.com)*. Mr. Malfatt is a Senior Engineer in the Worldwide Packaging and Test Products organization within the IBM Systems and Technology Group. He joined IBM in 1971 after receiving a B.S.E.E. degree from New Jersey Institute of Technology that same year. He also received an M.B.A. degree from Marist College. While at IBM he has held both semiconductor and packaging technical and management positions. His focus for the past 25 years has been specific to packaging, and he has written numerous papers on the subject. His current activities include interconnect development and technology limitations in high-current applications.

Daniel N. de Araujo *IBM Systems and Technology Group, 11400 Burnet Road, Austin, Texas 78758 (dearaujo@us.ibm.com)*. Mr. de Araujo has eight years of experience in board and chip design, simulation, and validation in high-end servers and high-volume consumer desktops. He joined the IBM Personal Systems Division in 1997 in Research Triangle Park, North Carolina. In 2001 he moved to Austin, Texas, and is currently a Senior Engineer in the eServer* xSeries Electrical Interconnect and Packaging Design group. He received a Master of Science degree in computer engineering from North Carolina State University in 2000 (*magna cum laude*) and a Bachelor of Science degree in electrical engineering from Michigan State University in 1997 (*summa cum laude*). Mr. de Araujo is a member of IEEE, Tau Beta Pi, and Eta Kappa Nu; he has one patent and 21 publications.

Stefano Oggioni *IBM Systems and Technology Group, Via Lecco 61, Vimercate, Italy (stefano@it.ibm.com)*. Mr. Oggioni is a Senior Engineer in the Worldwide Packaging organization of the IBM Systems and Technology Group. He joined IBM in 1981 after receiving his diploma in industrial electronics. Since then, he has held several positions in electronic card assembly and packaging development engineering, focusing on materials sciences and related industrial processes. Over the years he has had numerous foreign assignments and residencies at various IBM laboratories in the U.S. Mr. Oggioni has been involved in packaging development since 1989 and in advanced application designs, working primarily in the development of flip-chip processes on organic and flex carriers, since 1991. In 1995 he moved into the development of advanced laminate carriers and module assembly processes, leading the development of multichip modules (MCMs) and high-speed applications on organic substrates. Mr. Oggioni has received an IBM Outstanding Achievement Award and a Leadership Award. He has authored or co-authored several publications covering various aspects of electronic packaging. He holds eight U.S. patents and has achieved the IBM Sixth Level Invention Achievement Plateau; in 2003 he was named an IBM Master Inventor. Mr. Oggioni is a member of IMAPS.

Luke D. LaCroix *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (llacroix@us.ibm.com)*. Mr. LaCroix is a Senior Engineer in the IBM Systems and Technology Group. He joined IBM in 1974 and has worked in numerous areas including test and package development. He is currently in the ASIC package development area, focusing on image/package development and ASIC package strategies. Mr. LaCroix holds a patent and has numerous publications in his field of expertise.

Jamil A. Wakil *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (jwakil@us.ibm.com)*. Mr. Wakil currently works on package thermal development for IBM microelectronics, focusing on first-level thermal enhancement for organic packages. He holds an M.S. degree in mechanical engineering from the University of Texas at Austin, a B.S. degree in mechanical engineering from Texas A&M University, and a B.S. degree in electrical engineering from the University of Texas at Dallas. He has been with IBM for five years and has several patents and publications.

Gareth G. Hougham *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (ggh@us.ibm.com)*. Dr. Hougham is a Research Staff Member at the IBM Thomas J. Watson Research Center. He joined IBM in 1983 and has worked in the areas of polymer science and microelectronic packaging. He received his Ph.D. degree in polymer chemistry from Polytechnic University in 1992. His areas of interest have included synthesis and dielectric characterization of polyimides, fluoropolymers, negative-CTE materials, MEMs, and low-force LGA connectors. Dr. Hougham is past chairman of the Polymer Analysis Division of the Society of Plastics Engineers; he has authored or co-authored 18 issued patents, 35 publications, and a book on fluorine-containing polymers.

Nam H. Pham *IBM Systems and Technology Group, 11400 Burnet Road, Austin, Texas 78758 (npham@us.ibm.com)*. Mr. Pham received a B.S.E.E. degree from Michigan Technological University in 1988 and an M.S.E.E. degree from Syracuse University in 1991. He joined the IBM Microelectronics Division in East Fishkill, New York, in 1988 as a semiconductor reliability engineer. In 1993, he joined the IBM PowerPC* Somerset Design Center, Austin, Texas, where he worked on application and packaging development for PowerPC microprocessor designs. He is currently a Senior Engineer in the eServer xSeries Electrical Interconnect and Packaging Design group. He has one U.S. patent and has published 23 technical papers in various IEEE publications.

David J. Russell *IBM Systems and Technology Group, 1701 North Street, Endicott, New York 13760 (daviruss@us.ibm.com)*. Dr. Russell received his Ph.D. degree in chemistry and joined IBM Endicott in 1984. He has been in process development and has worked with materials and photoprocessing, including photoresist, solder mask, and photo-imageable dielectrics for build-up materials. While based in Endicott, he has supported operations in several IBM locations, including the IBM SLC operation in Yasu, Japan. His current position is in supplier technology development, working with build-up technology from several suppliers.